

Classifying Syntactic Regularities for Hundreds of Languages

Reed Coke
Department of EECS
University of Michigan
Ann Arbor
reedcoke@umich.edu

Ben King
Ubiquiti, Inc.
benking@ubiquiti.com

Dragomir Radev
Department of EECS
School of Information
University of Michigan
Ann Arbor
radev@umich.edu

Abstract

This paper presents a comparison of classification methods for linguistic typology for the purpose of expanding an extensive, but sparse language resource: the World Atlas of Language Structures (WALS) (Dryer and Haspelmath, 2013). We experimented with a variety of regression and nearest-neighbor methods for use in classification over a set of 325 languages and six syntactic rules drawn from WALS. To classify each rule, we consider the typological features of the other five rules; linguistic features extracted from a word-aligned Bible in each language; and genealogical features (genus and family) of each language. In general, we find that propagating the majority label among all languages of the same genus achieves the best accuracy in label prediction. Following this, a logistic regression model that combines typological and linguistic features offers the next best performance. Interestingly, this model actually outperforms the majority labels among all languages of the same family.

1 Introduction

Linguistic typology is a subfield of linguistics concerned with understanding the various patterns that are present across the world’s languages and how languages can be grouped via these patterns as well as other historic and geographic factors. Other fields of linguistics, such as historical linguistics and the study of endangered languages, depend heavily on knowledge drawn from typological comparisons. Linguistic typology is also useful in many areas of natural language processing.

WALS is an expansive linguistic resource useful in a variety of different NLP applications. WALS has data for over 2,500 languages regarding almost 200 rules compiled by a group of 55 authors spanning phonology, syntax, and lexicology. Examples of such rules and their possible values are given in section 3.2. WALS has been used as a lens through which to view worldwide typological relations (Lit-tauer et al., 2012) as well as a tool with which to discover which linguistic rules are rare across the world’s languages (Cysouw, 2011). Unsurprisingly, it has also served as a baseline for typological similarity measurements (Berzak et al., 2014). In this work the authors also show that the accuracy of the baseline increases with the number of WALS rules considered, but note the sparseness of these rules. Indeed, the sparseness of WALS has been noted by several researchers (Georgi et al., 2010; Cysouw and Comrie, 2008; Teh et al., 2008). In fact, currently the average rule in WALS has data for only 400 languages. This represents a matrix that is only 16% populated. To this end, we examine methods for automatically identifying rule values in the WALS database for the purpose of reducing sparseness in WALS. Our methods use the WALS database as of March 2015 as well as an additional copy of the Bible in each target language downloaded from `bible.com` or `bible.is` and word-aligned to an English version.

Despite its sparseness, WALS is still an excellent resource. Consider a similar resource, Syntactic Structures of the World’s Languages (SSWL) (Collins and Kayne, 2011). SSWL is considerably less sparse, with 56% of all possible data for 112

rules and 251 languages. In reality, many of these rules actually represent pieces of a single rule in WALS - SSWL rule 3 is *Verb Object* while SSWL rule 4 is *Object Verb*. In WALS, both of these rules are possible values for rule 83A: Order of Object and Verb. Thus, SSWL lacks the breadth that makes WALS such an appealing resource. WALS has an order of magnitude more languages and many more rules. In addition, WALS has genealogical information about each language that SSWL lacks.

While many of the rules in WALS, such as rule 83A: Order of Object and Verb, may seem simple, properties of linguistic universals can be exploited to extrapolate further information. Joseph Greenberg proposed a set of linguistic universals (Greenberg, 1963) that refer to statistical tendencies across the world’s languages. Further universals have been described since, particularly with respect to word order (Hawkins, 1983). While our system does not directly make use of these universals, the core mechanic of classifying rules based on the labels of other rules for the same language is central to our experimentation. This is the grounding theory behind the use of typological features.

Applying effective, automated methods for expanding this database, or similar databases, will not only create a better resource, but will also lead to improvements in multilingual NLP via better parsing, a better understanding of linguistic typology, and many other basic NLP tasks such as machine translation (Mikolov et al., 2013) and part-of-speech tagging (Das and Petrov, 2011). In addition, having a more expansive resource will allow for NLP to be leveraged in many under-resourced languages not only in Europe, but around the world.

The rest of this paper is structured as follows: Section 2 discusses other work in the field of NLP that has made use of features in WALS; Section 3 describes in detail how we generated our feature vectors and present the classification models we compared; Section 4 presents our results and provides an analysis; Section 5 concludes our work; and Section 6 suggests future work.

2 Related Work

Typological similarity has previously been shown to correlate with genealogical similarity both in the

fields of NLP (Rama and Kolachina, 2012) and historical linguistics (Dunn et al., 2005). Due to this, WALS has been used to study linguistic typology via computational methods. In order to determine language similarity via genealogical relatedness, it is important to know which rules are more telling of historical relatedness. To do this, (Wichmann and Kamholz, 2008) measured the variance of linguistic rules within language families, at the genus level, to determine their stability. Rules that change less often within language families, then, are more indicative of historical relationships when they are shared between languages. To do this, the authors extracted data for language families from WALS and calculated the probability of features being shared within and across language families. Feature variation was shown to be significantly different across genera. This indicates that language families are more similar internally than across genera, which in turn supports the concept of using typological similarity to predict language similarity.

These syntactic regularities are not only predictable within language genera, however. (Daumé and Campbell, 2007) used WALS as a database to discover implicative associations, similar to Greenberg’s universals. Not only did they recover many of Greenberg’s universals, but they also uncovered a host of other implicatures.

In similar work, (Cysouw and Comrie, 2008) sought to determine how consistent certain linguistic rules are in an effort to determine languages that are “more central for the structure of human language.” The authors cite “widely varying frequencies of available data” as a difficulty in their study, but conclude that many word order features seem to be central rules in language. Among these are rules 83A, 85A, 86A, 88A, and 107A, five of the six rules considered in this study.

Apart from typology, the data from WALS have also been used in natural language processing to advance several core NLP tasks. (Naseem et al., 2012) used language similarity information to improve multilingual parsing by defined a distance metric over WALS rules. The authors considered six WALS rules to create their similarity metric, three of which (85A, 86A, and 88A), are also used in this study.

(Bender et al., 2013) attempted a similar task, us-

ing interlinear glossed text (IGT) to predict typological features such as major constituent order and case system. (Lewis and Xia, 2008) also predict typological features from IGT by first learning a context free grammar for a language and then examining its structure. Both studies use data from WALS to evaluate their performance. We address a similar task, but we utilize projected dependency parses from English text as a source of knowledge for each foreign language, as many do not have extensive resources of their own. This is similar to IGT, but does not contain morphological information to the same extent. This method of projecting dependencies has been used to bootstrap linguistic resources in the past (Xia, 2007), (Hwa et al., 2005).

3 Experiments

We first created a corpus of Bibles word-aligned to an English Bible. The 325 languages represented in the corpus are very diverse, ranging from French, German, and Modern Standard Arabic to Acholi, Basque, and Tamil. Note that translated Bibles exist for many other languages that were not included in this study. In total, we use over 2 million aligned sentences. The rules were selected based on the frequency of their use in previous NLP research, a measure of usefulness. More information about them is provided in Section 3.2. For each rule we ran experiments using a variety of classifiers and feature vector combinations. All results are given in Section 4.

3.1 General Textual Feature Extraction

For each rule described in Section 3.2, we extract English dependencies according to certain criteria unique to that rule. We then project these dependencies onto the foreign language biblical sentences using word alignments determined by BerkeleyAligner (DeNero and Klein, 2007; Liang et al., 2006). Using these inferred dependencies, we calculate the feature vectors for the rule in question by the policy described in Section 3.2.

To give a simple example of projecting word alignments, consider the English sentence and its Ma'di translation in Figure 1. *is* and *resurrection* are the verb and object of the English sentence, while *i* and *onzika* are the Ma'di alignments. These alignments are shown using blue edges. As these words

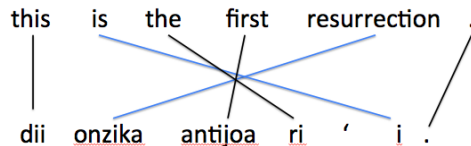


Figure 1: Word alignments for an example sentence in Ma'di and English

are aligned, we assume they serve the same grammatical function. We can see that the English sentence shows a verb-object ordering, while the Ma'di sentence demonstrates an object-verb ordering. This would then add one to the count of object-verb orderings to the linguistic feature vector of Ma'di for rule 83A.

Next, we give a detailed example of how the features for WALS rule 83A *Order of Object and Verb* can be drawn from projected dependency parses of English text as in the German example in Figures 2 and 3.

In this example, German would receive one count for having the same verb-object ordering in any phrase due to the (*Dothan, found*) dependency and one count for different ordering due to the (*brothers, after*) dependency. This sentence is not a question or a dependent clause, so the corresponding same/different counts would not be changed for those features. This process is then repeated for all aligned sentences in all languages with a known label for the rule.

3.2 Specific Feature Extraction

Each rule is described in detail below, along with definitions of its possible classes, rounded distributional information over the languages considered in this paper, and a brief motivation for why we chose to classify this rule. All parses referred to in this section were obtained using a Hash Kernel parser (Bohnet, 2010) and represented in the CONLL dependency parsing format. We restrict our system to instances where the English sentence aligns to only one foreign language sentence. The count of languages given for each rule is the size of the set of languages whose label for this rule is given in WALS and for whom we have a Bible.

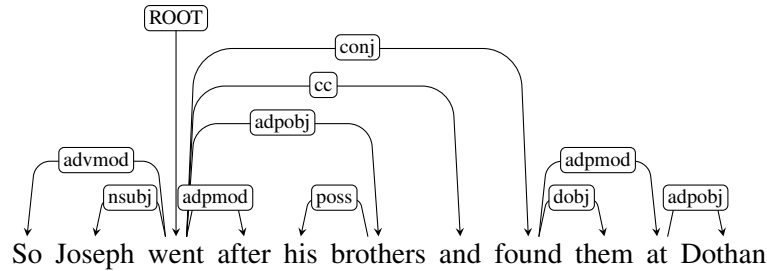


Figure 2: Sample English dependency parse

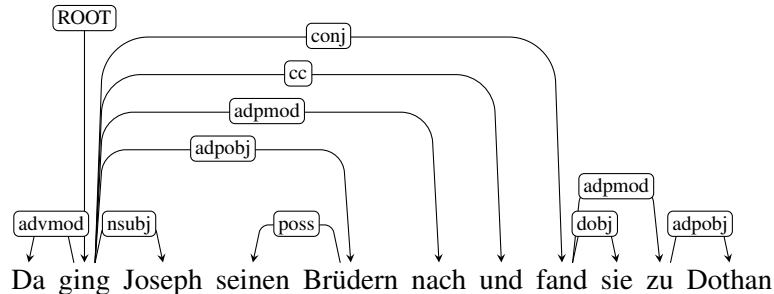


Figure 3: Sample German dependency parse

3.2.1 Rule 83A: Order of Object and Verb

Rule 83A defines the order in which the direct object and verb of a sentence occur. Its label is known for 299 languages. The three possible classes are *Object Verb* (The man saw the dog. 38%), *Verb Object* (The man the dog saw. 58%), and *No Dominant Order* (both valid in certain conditions. 4%). Besides being fundamental to any sort of parsing or chunking task, many other typological rules have been found to be statistically related to the order of subject, object, and verb (Greenberg, 1963). As it has also been found that subject generally precedes object (Greenberg, 1963), knowing the order of the object and verb allows not only accurate guesses about the order of subject, object, and verb, but also about many other correlated rules. To identify all occurrences, we examine all dependencies with the label *dobj*. The instance is kept as long as both of the English and foreign object words align only to each other and both of the English and foreign verb words align only to each other. The resulting feature vector has six columns: the count of all sentences in which the verb-object ordering is identical between English and the foreign language; the count of all sentence in which the ordering is reversed; the count

of all questions in which the ordering is identical; the count of all questions in which it is reversed; and finally the count of dependent clauses where the order is identical and the count of dependent clauses in which the order is reversed.

3.2.2 Rule 85A: Order of Adposition and Noun Phrase

Rule 85A defines the order in which adpositions occur relative to their governing noun phrase. Its label is known for 256 languages. The five possible classes are *postpositions* (The dog went the park to. 43%), *prepositions* (The dog went to the park. 49%), *inpositions* (The dog went the to park. 0%), *more than one type with none dominant* (multiple valid. 6%), and *no adpositions* (The dog went the park. 1%). This rule is critical for having any sort of successful semantic parse or narrative processing. We examine all dependencies with the label *adpmod* and have a parent word in the sentence. The instance is kept as long as the English and foreign adpositions align only to each other and the English and foreign governing nouns align only to each other. The feature vector has a dimensionality of six and consists of the same features as the feature vector for Rule 83A, simply computed for the order of the adposi-

tion and noun phrase, rather than object and verb.

3.2.3 Rule 86A: Order of Genitive and Noun

Rule 86A defines the order in which genitives appear with respect to their governing noun. 254 languages in our dataset have a known label for this rule. The three possible classes are *Genitive-noun* (My dog 47%), *Noun-genitive* (Dog my 45%), and *Both occur, none dominant* (Both valid 8%). This rule is absolutely necessary for coreference resolution. We examine all dependencies with the label *poss*, as long as the genitive and noun align only to their respective foreign language counterpart. Again, we use a length six feature vector with the same procedure over this set of dependencies.

3.2.4 Rule 88A: Order of Demonstrative and Noun

Rule 88A defines the order in which demonstratives appear together with nouns. There are 171 languages in our set of Bibles with a known label for this rule. The six possible classes are *Demonstrative noun* (This dog is shaggy. 50%), *Noun demonstrative* (Dog this is shaggy. 42%), *prefix on noun* (This-dog is shaggy. 0%), *suffix on noun* (Dog-this is shaggy. 2%), *Demonstrative noun demonstrative* (This dog this is shaggy. 2%), and *two or more of the previous options, none dominant* (multiple valid. 4%). This rule is also necessary for coreference resolution. We select all dependencies that are of type *det* or *pron*, have a lexical value of *this*, *that*, *these*, or *those*, and whose parent is a noun. We compute the feature vector in the same fashion as the above rules.

3.2.5 Rule 92A: Position of Polar Question Particles

Rule 92A defines the position that a polar question particle appears in a sentence. Its rule is known for 91 languages. Polar question particles signal grammatically that the sentence is a yes or no question. The six possible classes are *beginning of sentence* (23%), *end of sentence* (36%), *second word in sentence* (3%), *anywhere else* (2%), *either of two positions* (1%), *no question particle* (34%). In languages with polar questions, this rule would allow for much higher quality question answering and dialogue systems. Our selection of this rule also

demonstrates that the phenomenon does not even need to occur in the source language in order to be approached. We select all dependencies that contain questions. We then use the presence of an English *wh-* word to sort these questions into polar and information questions. Using these labels, we examine the foreign language questions for the word that appears most often in polar questions but not information questions using the ratio of relative frequencies. We then count the number of times this word appears in each of the possible positions. The feature vector has a length of four, and each entry corresponds directly to a label class: we count how often the inferred question word appears initially, second, last, or elsewhere.

3.2.6 Rule 107A: Passive Constructions

107A describes the presence or absence of passive constructions. We have Bibles for 93 of the WALS languages with a known label for this rule. The two possible classes are *passive construction* (The dog was seen by the butler. 53%) and *no passive construction* (The butler saw the dog. 47%). This rule was chosen not only because it would allow for better grammar induction, a very relevant task for such under-resourced languages, but also to demonstrate that our system and approach can handle rules other than simple word order rules. We select all dependencies that are of type *nsubjpass*, signaling a passive subject, and *nsubj*, signaling an active subject. We then determine the number of times that the order of subject and verb differs between these two sentence types in the foreign language. The feature vector is built identically to all other rules except for 92A.

3.3 Text Classifier Training

In order to train the text classifiers, we used word-aligned biblical texts in English and each of the languages considered along with the labels from WALS. For each rule discussed in Section 3.2, we select dependencies and create normalized feature vectors according to the discussed method. This forms the Text feature set in Section 4.

3.4 Typological Classifier Training

We also consider that some labels can be predicted from the labels of other rules within a language. We

create feature vectors for each language using five of the six rules discussed in Section 3.2 in order to classify the sixth rule. This comprises the Rules feature set in Section 4.

3.5 Genealogical Classifier Training

Previous work has shown that propagating knowledge from genealogically similar languages has demonstrated markedly better results than propagating knowledge from a random language (Naseem et al., 2012). For our purposes we consider the genus and family of each language as given by WALS. To this end, we consider simply propagating the majority label from all languages of the same genus and, separately, all languages of the same family. We also create feature vectors such as those for the Text and Rules features, as described above.

4 Results and Discussion

Majority is simply the majority class as shown in 3.2. Naïve Bayes is run using the default implementation in Weka (Hall et al., 2009). For logistic regression, as we do not have a developmental data set, we simply experiment with five common regularization parameters - 1.0, 0.5, 0.1, 0.01, and 10^{-8} (the default value in Weka). As we are only evaluating over languages that already have known labels in WALS, we consider the accuracy of each classifier through leave-one-out cross-validation.

4.1 Simple Linguistic Features

It is clear from Tables 1 - 6 that the combination of textual and typological features works best. It can also be seen that almost regardless of the regularization parameter, logistic regression models this data with higher accuracy than naïve Bayes. In the cases that this is not true, rule 83A and rule 88A, it is the linguistic features that perform best. This is also encouraging, as the main problem we are trying to address is the sparseness of the WALS data.

There are also two rules that perform quite poorly, rule 92A and rule 107A, depicted in Tables 5 and 6. These rules are significantly more subtle linguistically than the rest, which are relatively straightforward for our parser to detect accurately. The complexity of the pipeline for rule 92A and the numerous ways in which languages can mark passive con-

structions, regarding rule 107A, are simply difficult to perceive.

We also find that regularization does make a difference for the various logistic regression models. In some cases, such as Table 5, the difference can be as large as 10% accuracy. In others, such as Table 3, the difference is only 0.8%. In general, it appears that a regularization parameter between 1 and 0.5 yields the best, most consistent results for the combined feature vector case. This is likely due to the nature of the two feature vectors being combined, as one is inherently discrete while the other is continuous, but more experimentation is necessary to determine the relationship more precisely.

4.2 Genealogy-based Classification

The genealogical experiments led to rather different results. We first tested a very simple approach, assigning the majority label for languages of the same genus or same family. Table 7 demonstrates these results.

It is very clear that this majority voting scheme works much better than the other approaches. This is particularly interesting, given the argument that typological features have been claimed to be an effective proxy for genealogical data (Rama and Kolachina, 2012).

We also attempted to incorporate the genealogical features as a third feature source, similar to our textual and typological features. Combining genealogical features with the either Text or Rule features did not outperform the systems already presented, nor did the combination of all three for the naïve Bayes or any logistic regression model. We also considered that perhaps there is something intrinsic to majority voting that is boosting the performance of this simple approach. To confirm, we ran experiments on the Text and Rule feature vectors using k -nearest neighbor classification. Again, these experiments did not rival the performance of our other classifiers, regardless of the choice of k .

It has been estimated that the accuracy of the current WALS database is possibly close to 96%, based on examination of the entries for Latvian (Cysouw, 2011). The results in Table 7 are just shy of this number for many of the rules. This means that for many rules, the simple genealogical majority measure could be used as is to expand WALS. Obviously

Features	Majority	Naïve Bayes	LR1	LR.5	LR.1	LR.01	LR-8
Text	58.0%	92.0%	92.3%	92.3%	92.3%	92.3%	92.0%
Rules	58.0%	82.2%	83.3%	83.3%	83.6%	83.6%	83.6%
Text+Rules	58.0%	90.6%	90.6%	90.0%	88.9%	89.6%	89.0%

Table 1: Accuracy for rule 83A: Order of Object and Verb N=299

Features	Majority	Naïve Bayes	LR1	LR.5	LR.1	LR.01	LR-8
Text	49.6%	82.0%	83.2%	83.2%	84.4%	84.4%	84.0%
Rules	49.6%	82.0%	84.8%	84.8%	84.4%	84.0%	83.6%
Text+Rules	49.6%	83.2%	86.7%	85.5%	84.8	83.2%	82.4%

Table 2: Accuracy for rule 85A: Order of Adposition and Noun Phrase N=256

Features	Majority	Naïve Bayes	LR1	LR.5	LR.1	LR.01	LR-8
Text	47.2%	66.9%	66.5%	66.9%	66.1%	66.1%	66.1%
Rules	47.2%	81.6%	81.3%	80.5%	80.5%	80.1%	79.8%
Text+Rules	47.2%	79.1%	81.1%	81.5%	81.9%	81.5%	81.1%

Table 3: Accuracy for rule 86A: Order of Genitive and Noun N=254

Features	Majority	Naïve Bayes	LR1	LR.5	LR.1	LR.01	LR-8
Text	50.2%	74.3%	83.6%	83.6%	84.2%	84.2%	84.2%
Rules	50.2%	60.8%	67.3%	66.7%	64.9%	64.3%	63.7%
Text+Rules	50.2%	74.2%	81.3%	81.3%	80.7%	78.9%	76.0%

Table 4: Accuracy for rule 88A: Order of Demonstrative and Noun N=171

Features	Majority	Naïve Bayes	LR1	LR.5	LR.1	LR.01	LR-8
Text	36.2%	26.4%	27.5%	27.5%	29.7%	28.6%	27.5%
Rules	36.2%	38.5%	37.4%	36.3%	40.7%	39.6%	38.5%
Text+Rules	36.2%	33.0%	42.9%	35.2%	34.1%	33.0%	31.9%

Table 5: Accuracy for rule 92A: Position of Polar Question Particles N=91

Features	Majority	Naïve Bayes	LR1	LR.5	LR.1	LR.01	LR-8
Text	52.6%	50.5%	47.3%	48.4%	47.3%	45.2%	44.1%
Rules	52.6%	55.9%	55.9%	58.1%	58.1%	58.1%	55.9%
Text+Rules	52.6%	57.0%	58.1%	59.1%	57.0	55.9%	54.8%

Table 6: Accuracy for rule 107A: Passive Constructions N=93

this system will not be as accurate as a trained linguist, and we are not suggesting that if we expanded WALS this way it would be fit to replace the current database, but rather that the system could be used to create a more expansive database at the cost of slightly lower accuracy.

5 Conclusion

In this study, we examined the effectiveness of classifying a series of syntactic rules across a set of foreign languages. This is doable by considering the values of other syntactic rules in the database. However, it is also worth considering actual linguistic features for this linguistic task, as well as combin-

Rule	Best Other	Same Genus	Same Family
83A	92.3%	95.8%	86.3%
85A	86.7%	94.3%	84.0%
86A	81.9%	93.4%	80.8%
88A	84.2%	88.0%	81.3%
92A	42.9%	91.1%	73.3%
107A	59.1%	95.9%	81.3%

Table 7: Accuracy using Genealogical Neighbors

ing the two disparate sources. Finally, we also considered genealogical data available from the WALS database.

We have shown that with proper regularization, we can achieve better classification accuracy by augmenting the currently available WALS data with linguistic features. Though in some cases the individual linguistic features actually outperform the combined features, this is not generally the case and it only applies when the linguistic features are performing well on their own. In addition, it is not possible to say from our results which regularization parameter is best. We do, however, demonstrate that this combination of features can lead to greater success than either feature independently. This result is important despite performing worse than the genealogical approach as the goal of our study is ultimately to find accurate methods to classify features for languages that are not already contained in WALS as well as features that are missing for languages already present.

Considering only the genus of the language, we are within 2% of achieving human-like accuracy for half of our six rules. This is encouraging and further experimentation needs to be done over the entire set of rules in WALS. However, this approach is limited as the fundamental problem of sparsity applies to typological data, which is also used to determine the true testing labels in this case. For this reason, it is also important that we have demonstrated the usefulness of purely linguistic features that are separate from the WALS database, as well as their compatibility with the typological values in WALS.

6 Future Work

Moving ahead, more complex WALS rules could be addressed with this method. While the rules we

consider are all useful in NLP research, there are other rules that are almost as frequently used that are more complex, such as those governing negation. In particular, modules concerning morphological features that are not easily detectable using word alignments are the next logical direction. It would also be worth gathering word aligned Bibles from other well-supported source languages besides English. As was demonstrated by Rule 92A: Position of Polar Question Particle, the system can work even when the feature does not appear in the source language, but it is certainly more difficult to account for. Having several source languages would minimize the number of rules for which this is the case.

In addition, experiments with combinations of genealogical, textual, and typological features should be performed. However, rather than simply concatenating these features, the features themselves should be combined into compound features, such as "Germanic and 83A-No_Dominant_Order". We believe these types of features could improve upon the performance of the majority-vote features that are currently the best measure. The system might then achieve even higher accuracy, possibly on par with human performance for some rules, allowing the WALS database to be automatically expanded, perhaps as a parallel resource to the current database.

In addition, having shown that this is a reasonably accurate method for determining values of WALS rules, the method should be applied on as of yet unlabeled languages in the WALS database. We also have nearly 400 such Bibles in languages not contained in WALS, which have not been used in this experiment. Adding them to WALS would expand the number of languages in the database by more than ten percent. While we would not be able to apply the most accurate, genealogical method, we

could certainly make use of linguistic features from the Bibles.

It would also be possible to use semi-supervised methods such as semi-supervised k-Nearest Neighbors and graph regularization (Zhu et al., 2003) as an improvement to our classification method. Though we did experiment with k-Nearest Neighbors in this study, it is possible that a different approach could still yield good results. The idea of nearest-neighbors classification stands to be very useful, as shown by the majority label propagation. Graph regularization has been shown to be effective in weakly-supervised situations where there is sparse training data (Hassan et al., 2014). Therefore, it would be a good idea to apply it to the WALS data. One could build a similarity network of all the languages using the known WALS rules that they share in common and regularize this graph in order to propagate labels to some extent whenever possible. Constraints such as the genus-majority relationship discussed in this paper as well as constraints borrowed from linguistic universals, representing conditional dependencies between rules, could also be used to help regularize the entire network.

References

- [Bender et al.2013] Emily M. Bender, Michael Wayne Goodman, Joshua Crowgey, and Fei Xia. 2013. Towards creating precision grammars from interlinear glossed text: Inferring large-scale typological properties. In *Proceedings of the 7th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 74–83, Sofia, Bulgaria, August. Association for Computational Linguistics.
- [Berzak et al.2014] Yevgeni Berzak, Roi Reichart, and Boris Katz. 2014. Reconstructing native language typology from foreign language usage. *CoRR*, abs/1404.6312.
- [Bohnet2010] Bernd Bohnet. 2010. Very high accuracy and fast dependency parsing is not a contradiction. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 89–97. Association for Computational Linguistics.
- [Collins and Kayne2011] Chris Collins and Richard Kayne. 2011. Syntactic structures of the world’s languages. <http://sswl.railsplayground.net>.
- [Cysouw and Comrie2008] Michael Cysouw and Bernard Comrie. 2008. How varied typologically are the languages of africa? In R Botha and C Knight, editors, *The Cradle of Language*, Studies in the Evolution of Language, pages 190–203. Oxford University Press, Oxford, UK.
- [Cysouw2011] Michael Cysouw. 2011. Quantitative explorations of the world-wide distribution of rare characteristics, or: the exceptionality of northwestern european languages. *Expecting the Unexpected*, pages 411–431.
- [Das and Petrov2011] Dipanjan Das and Slav Petrov. 2011. Unsupervised part-of-speech tagging with bilingual graph-based projections. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT ’11, pages 600–609, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Daumé and Campbell2007] Hal III Daumé and Lyle Campbell. 2007. A bayesian model for discovering typological implications. *CoRR*, abs/0907.0785.
- [DeNero and Klein2007] John DeNero and Dan Klein. 2007. Tailoring word alignments to syntactic machine translation. In *ANNUAL MEETING-ASSOCIATION FOR COMPUTATIONAL LINGUISTICS*, volume 45, page 17. Citeseer.
- [Dryer and Haspelmath2013] Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- [Dunn et al.2005] Michael Dunn, Angela Terrill, Ger Reesink, Robert A Foley, and Stephen C Levinson. 2005. Structural phylogenetics and the reconstruction of ancient language history. *Science*, 309(5743):2072–2075.
- [Georgi et al.2010] Ryan Georgi, Fei Xia, and William Lewis. 2010. Comparing language similarity across genetic and typologically-based groupings. In *Proceedings of the 23rd International Conference on Computational Linguistics, COLING ’10*, pages 385–393, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Greenberg1963] Joseph H. Greenberg. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In Joseph H. Greenberg, editor, *Universals of Human Language*, pages 73–113. MIT Press, Cambridge, Mass.
- [Hall et al.2009] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The weka data mining software: An update. *SIGKDD Explor. Newsl.*, 11(1):10–18, November.
- [Hassan et al.2014] Ahmed Hassan, Amjad Abu-Jbara, Wanchen Lu, and Dragomir R. Radev. 2014. A random walk based model for identifying semantic orientation. *Computational Linguistics*, 40.

- [Hawkins1983] John A Hawkins. 1983. Word order universals: Quantitative analyses of linguistic structure.
- [Hwa et al.2005] Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering*, 11:11–311.
- [Lewis and Xia2008] William D Lewis and Fei Xia. 2008. Automatically identifying computationally relevant typological features. In *IJCNLP*, pages 685–690.
- [Liang et al.2006] Percy Liang, Ben Taskar, and Dan Klein. 2006. Alignment by agreement. In *Proceedings of the Main Conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, HLT-NAACL '06, pages 104–111, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Littauer et al.2012] Richard Littauer, Rory Turnbull, and Alexis Palmer. 2012. Visualising typological relationships: Plotting wals with heat maps. In *Proceedings of the EACL 2012 Joint Workshop of LINGVIS & UNCLH*, EACL 2012, pages 30–34, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Mikolov et al.2013] Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *CoRR*, abs/1309.4168.
- [Naseem et al.2012] Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. Selective sharing for multilingual dependency parsing. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers - Volume 1*, ACL '12, pages 629–637, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Rama and Kolachina2012] Taraka Rama and Prasanth Kolachina. 2012. How good are typological distances for determining genealogical relationships among languages? In *Proceedings of the 24th International Conference on Computational Linguistics*.
- [Teh et al.2008] Yee Whye Teh, Hal Daum Iii, and Daniel Roy. 2008. Bayesian agglomerative clustering with coalescents. In *In Advances in Neural Information Processing Systems*.
- [Wichmann and Kamholz2008] Søren Wichmann and David Kamholz. 2008. A stability metric for typological features. *STUF-Language Typology and Universals Sprachtypologie und Universalienforschung*, 61(3):251–262.
- [Xia2007] Fei Xia. 2007. Multilingual structural projection across interlinear text. In *In Proc. of the Conference on Human Language Technologies (HLT/NAACL 2007)*, pages 452–459.
- [Zhu et al.2003] Xiaojin Zhu, Zoubin Ghahramani, John Lafferty, et al. 2003. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*, volume 3, pages 912–919.